Competence and Advice *

Anna Denisenko† Catherine Hafer‡ Dimitri Landa§

Oct 13, 2023

Abstract

We develop a theory of policy advice that focuses on the relationship between the competence of the advisor (e.g., an expert bureaucracy) and the quality of advice that the leader may expect. We describe important tensions between these features present in a wide class of substantively important circumstances. These tensions point to the presence of a trade-off between receiving advice more often and receiving more informative advice. The optimal realization of this trade-off for the leader sometimes induces her to prefer advisors of limited competence – a preference that, we show, is robust under different informational assumptions. We consider how institutional tools available to leaders affect preferences for advisor competence and the quality of advice they may expect to receive in equilibrium.

Introduction

Government leaders are inevitably defined by the outcomes of policies they choose, yet those policies are rarely a product of the leaders' *sui generis* judgments. Instead, leaders must often rely on advice from government officials better positioned to develop expertise and competence with respect to the questions of interest. It is tempting to assert, in light of this, that the leaders' political and policy successes owe a lot to the competence of their advisors.

*We thank Ethan Bueno de Mesquita, Wiola Dziuda, Sandy Gordon, Federica Izzo, Carlo Prato, Matthew Stephenson, and participants at the PECO 2022, Kellogg Political Economy Conference, and the NYU PE Workshop for valuable feedback on earlier versions of this paper.

†Ph.D. student, Wilf Family Department of Politics, NYU, e-mail: ad4205@nyu.edu

‡Associate Professor, Wilf Family Department of Politics, NYU, e-mail: catherine.hafer@nyu.edu

§Professor, Wilf Family Department of Politics, NYU, e-mail: dimitri.landa@nyu.edu

But this seemingly obvious claim obscures fundamental strategic complexities at the core of the leader-advisor relationship.

We develop a theory of policy advice that focuses on the connection between central elements of that relationship – advisor competence (agency expertise) and the quality and quantity of advice that the leader may expect – and calls into question some strongly held intuitions about the quality of advice and about the politics of policy-making more broadly.

Our theory highlights an underappreciated tension between advisor (agency) competence and the informativeness of advice in equilibrium: while more competent advisors sometimes give more informative advice, in a wide class of substantively important circumstances we characterize, we should expect the opposite to be the case. More broadly, our theory sheds light on when the trade-off between competence and information, or, equivalently, between better information and more information revelation, is likely to be acute and on what institutional features can mitigate it.

The policy-making setting to which our theory most readily applies is the interaction between political leaders and bureaucrats. In such settings, the leader might be an elected official or a politically appointed head of an established government agency, and the informed advisors are the agency's career civil servants. The following, intentionally somewhat stylized, examples from such settings evoke recent events and help motivate our analysis.

1. The political appointees of a U.S. President are seeking specific information about corrupt government practices in foreign country X, and the U.S. Ambassador to X, who is a recognized expert on business and political practices in X, is in a particularly good position to have such information. Furthermore, because he is understood to be exceptionally capable and a person of integrity, the information he provides would be regarded as reliable. The ambassador sees the long-term goal of his work to be supporting democratic institutions and political development in X, which he believes to be central to the U.S. interests in the region and to be best achieved by strengthening the rule of law in X. He also suspects that the political leadership he advises places

greater weight on the President's own immediate electoral fortunes than on the success of democratic practices in X, and that the leadership's sudden intense interest in corruption in X is borne of a hope of finding a pressure point against a political rival. This advisor, then, chooses to "play dumb" and risk dismissal rather than reveal to the leader what he knows.

2. The U.S. Consumer Finance Protection Bureau (CFPB) has a staff consisting mainly of civil servants attracted to the mission of "protect[ing] consumers from unfair, deceptive, or abusive practices" from firms in the financial services sector. The CFPB Director is a political appointee, whose preferences reflect the political interests and policy commitments of the sitting president. Suppose that the President and the Director are committed to minimizing government regulation of business and to promoting the profitability of big banks. The CFPB civil servants have been crafting specific rules for companies offering consumer credit; they are to offer policy options to the Director, but the Director has the power to choose which option to implement. Suppose the Director (the leader) wishes to choose the option that will be most advantageous to a subset of banks offering consumer credit, and that the CFPB staff (the advisor) has information on point but does not wish to reveal it because it believes that that should not be a relevant consideration. Instead, it provides information on other features of policy, and claims to be unable to obtain the information the leader desires.

3. The head of country Y's government strongly supports restarting the program of targeted assassinations of known leaders of terrorist organizations, which he values, in part, because they give him a chance to look tough and effective in protecting the country. Suppose that Y's head of an intelligence agency, on the other hand, has come to believe that targeted killings of terrorist leaders provide relatively little benefit in disrupting terrorist activity and distract the agency from its mission of developing intelligence about potential terrorist threats. While the leader presses the head of the

intelligence agency (the advisor) to identify the location of a desired target, the advisor may prefer to forgo gathering the relevant information or to fail to reveal it, even at the risk of angering the leader.

In these and many other instances, civil servants have preferences that suggest an interest in maintaining the status quo. Lawrence J. Peter (of the Peter Principle) puts this in a characteristically flip fashion: "Bureaucracy defends the status quo long past the time when the quo has lost its status." But while Peter's comment implies a certain degree of criticism, Huq and Ginsburg (2018) view this conservatism as a guarantor of political stability, and describe as essentially symmetric with respect to possible challenges to status quo: "Just as bureaucracy may make progressive reform difficult to achieve, it also slows down rapid shifts away from liberal democratic norms" (p. 129). The bureaucracies' status-quo bias need not imply their primitive conservatism: As a rule, bureaucrats have long time horizons and so are apt to discount the value of responding to what may be short-term trends. In contrast, the time horizons of elected officials tend to be short, and they prefer to match the policy to the immediate trends, pandering to the voters' prior beliefs or otherwise bolstering their electoral fortunes while discounting longer-term implications of their actions.

In our model, an advisor who, like the bureaucrats in the above description, prefers maintaining the status-quo policy and has access to superior information, chooses whether to reveal to the leader her privately observed signal about the state of the world. Assuming the advisor shares her signal, the leader updates more strongly if the advisor is commonly known to be more competent and, conditional on such an update, shifts policy farther. Given this expectation and the gap between the most preferred policies of the leader and of her advisor, a more competent advisor stands to lose more from revealing her signal than does a less competent one. In general circumstances, the less competent advisors will, thus, have stronger incentives to reveal their information to the leader than will more competent ones. The leader, then, faces a trade-off between the quality of advice he receives and the likelihood of receiving it.

We describe conditions that influence the advisors' incentives to share and obtain information, as well as conditions that determine when leaders may be better off having advisors with lower competence, even though it means that the advice they receive is less reliable. We also show that the requirement of disclosure of advisors' information (a common consideration in the context of relationships between political leaders and bureaucratic agencies) has an equivocal effect on the leader's utility. When the advisor's information is fixed, it allows the leader to extract more information, and is more beneficial when the advising agency is more competent. But the opposite is true when the advisor's information is endogenous: requiring disclosure encourages more competent advisors not to acquire information, and the leader should limit requiring disclosure to when the advisor is relatively incompetent – that is, when the value of information that the enforced disclosure could reveal is, ultimately, not too great.

The rest of the paper proceeds as follows. Following a brief review of the prior literature, we introduce our model. We then begin its analysis with a general formulation of the trade-off between advisor competence (i.e., the quality of the information available to the advisor) and the strategically supportable quantity of advice offered, and then study the conditions that generate this trade-off in equilibrium. After that, we extend the model to include endogenous information acquisition, modeled as an experiment with non-contractible precision, and let the Advisor choose both whether to conduct the experiment and the experiment's precision. We show, first, that the only precision consistent with robust equilibrium play is the maximum possible precision, and that the Advisor's competence affects her incentives to obtain information in a way that echoes her incentives to reveal information in the fixed-information environment, but that such incentives are moderated by the publicity of the experiment. In the last substantive section, we introduce the possibility of requiring disclosure of Advisor's information and study its effects on the Advisors' incentives to acquire information and the Leader's equilibrium utility.

Connection to the Literature

The relationship between leaders and their advisors is critically affected by three sources of agency problems: (1) advisors have their own policy preferences; (2) advisors have informational advantage over leaders; and (3) advisors may have differing abilities to obtain information, which affects the quality of advice they could give and choose to give to the leaders. The primary focus of the previous work has been the relationship between the first two of these three factors. We focus, instead, on how the quality of an advisor's information affects the leader-advisor relationship.

Information revelation through advice has been the focus of a substantial literature in political economy, including in the political economy of bureaucracy. For a review focusing on the themes of delegation and communication within hierarchies, see Gailmard and Patty (2012). One branch of this literature models communication as "cheap talk" in which advisors' set of feasible messages is independent of their information. Early influential studies in this branch include Gilligan and Krehbiel (1989), Austen-Smith (1990), and Austen-Smith (1993). (For an important review of this literature, see Sobel (2013).) A key result in this literature is that the divergence in the actors' preferences limits revelation, and successful communication requires that the advisor's and the leader's preferences be sufficiently aligned. (Important exceptions are Battaglini (2002) and Aybas and Callander (2023), who identify policy-relevant settings in which communication fully favors, respectively, the receiver, and the sender.)

In contrast with this literature, this paper models the communication of hard evidence, rather than cheap talk, and focuses not on the effects of preference heterogeneity, but, holding the difference of the actors' preferences fixed, on the effects of the advisor's competence on her incentives to reveal existing, and to obtain additional, information. We model the Advisor's communication as verifiable messaging for both pragmatic and substantive reasons. To understand the pragmatic reason, first note that if the advisors can distort the signal in messaging the leader, then in the setting we study, advisors of all types will have an incentive to do so maximally – that is, sending the messages that are maximally uninformative. Given our interest in understanding the implications of the variation in advisor competence for the quality of their advice to the leader, unverifiable messages are, thus, not the right informational setting for the analysis. In contrast, in the context of verifiable messages, the variation in advisor type induces important differences in the nature of communication, suggesting that that is an appropriate setting for our analysis. The second set of reasons is substantive. To begin with, the evidence available to the bureaucratic agencies is typically in the form of internal reports, distorting which would create substantial administrative and legal jeopardy. And further, the leader's policy action based on the received advice must often engage directly with the evidentiary basis of the message (e.g., verifying the presence and location of an assassination target before and after the hit, or showing the evidence of corruption in an attempt to coerce a foreign leader to take a particular action). Given such constraints, a more natural way to think of the possibility of distortion is with respect to the (im-)precision of the information that the agency chooses to acquire, and we study such a possibility in an extension of our baseline model below.

Another branch of the literature on strategic communication focuses precisely on this kind of hard information. A touchstone finding in this literature is that all private information is revealed in equilibrium (Milgrom, 1981, 2008), and a major focus of subsequent work has been on understanding when the unraveling logic of the full revelation result does not hold. An important early paper is Shin (1994), which shows that the leader will not be able to infer perfectly the advisor's private information in the event of "no news" when the advisor's knowledge is imperfect; see also Wolinsky (2003). Dziuda (2011) shows that, in a setting where the fixed expert's preferences are different and unknown to the decision-maker, there is never full disclosure, but the expert offers pros and cons for the advocated alternative in order to pool with the honest/non-strategic type. Che and Kartik (2009) show that, when information acquisition is endogenous and the advisor shares the leader's preferences, the latter prefers an advisor whose prior beliefs are different than her own because it incentivizes the advisor to acquire information in hopes of persuading the leader.

A seminal paper on public information acquisition by an interested agent is Kamenica and Gentzkow (2011). Unlike the setting in that paper and in the subsequent extensive literature, the public experiment setting we study is one in which the precision of the experiment is not (directly) observable. Di Pei (2015)) and Gentzkow and Kamenica (2017) study cheap-talk and verifiable-information models, respectively, in which the agent acquires information covertly and then decides whether to reveal it (in the latter paper, with the relevant received evidence). A key result in these papers is that requiring disclosure does not alter the equilibrium play. As indicated above, this result does not hold in our model. Because explicating the reasons for this contrast builds on the detailed analysis below, we defer the discussion of this contrast until later in the paper, in the section on disclosure requirement.

A key rationale behind the downside of transparency explored in the literature is that transparency encourages the agents to pander to their principal's prior beliefs (in political settings, see, e.g., Canes-Wrone, Herron and Shotts (2001); Fox (2007); Patty and Turner (2021); and Turner (2022)). The mechanism that underlies the negative consequences of the disclosure requirement in our model is different – it comes from the effect of discouraging the information acquisition by the agent due to information leakage and is related to the "stovepiping" studied in Gailmard and Patty (2013).[1]

Bhattacharya and Mukherjee (2013) find that improving an advisor's quality in a panel of experts can lower the decision-maker's utility. A necessary condition for such an outcome in their model is that the leader's default policy in the absence of revelation be sensitive to advisor quality. In contrast, in our model, the optimal default policy is constant in

[1]Somewhat more distantly, our model is also related to the literature on the effects and value of biased information. In the present study, we abstract away from the possibility of ideological bias on the part of either principals or agents (they are explored in detail in a companion paper.) In a seminal paper in that literature, Calvert (1985) shows that the policy-maker may prefer a biased adviser when the policy-maker is concerned about the extreme utility realizations from policy choices. Patty (2009), who examines the incentives stemming from the supply side of the advice, finds that a policy-maker's bias leads to more biased information collection by the advisor and a bias-reinforcing advice.

advisor quality, and the impact of the latter is, rather, channelled through the quality of information provided, allowing us to get a sharper characterization of the conditions under which the leader may be better off with a lower-quality advisor and to examine more fully its implications.

Unlike the above studies, our focus is on the strategic implications of the relationship between the advisors' informational advantage over leaders and the quality of advice they could give to the leaders. A standard intuition, captured in a number of political economy models, is based on the career-concerns rationale: an advisor wants to appear well-informed (competent). See Ottaviani and Sørensen (2006) for an explicit analysis of the effects of this motivation. In contrast, in our model, the advisor's incentives are such that being well-informed can make her less useful to the leader and lower her own utility.

Several models focusing on the process of information/expertise acquisition show that there is an upside to starting with relatively less competent agents. Sobel (1993) shows that when the agent is uninformed, she might exert more effort to achieve an outcome than would a better informed agent, leading to the possibility of the former being more attractive to the principals. Gailmard and Patty (2007) model bureaucratic competence as an exogenous cost of acquiring expertise and show that when the cost is too high, the legislature does not reward the advisor for expertise acquisition.

The Baseline Environment

We analyze a strategic interaction between a Leader (he) and an Advisor (she) of known competence θ. The Leader wishes to choose an action that will match a state of the world, which he does not directly observe. Instead, the Leader may be able to obtain information about the state from his Advisor, whose competence determines the informativeness of the signal about the state of the world that the Advisor privately observes. The advisor's messages are verifiable; for the reasons discussed above, this is the natural environment in which to study the implications of advisor competence for the quality of advice to the leader.[2] The timeline of the game is as follows:

1. Nature determines the state of the world $w \in \mathbf{R}$, where w is a draw from a standard normal distribution $N(0, 1)$.

2. The Advisor of known competence θ observes signal s about the state of the world w, $s = w + \varepsilon$. The variable ε represents random noise drawn from a normal distribution with mean 0 and precision θ (s.t. $\theta \in \mathbf{R}^+$), $\varepsilon \sim N(0, 1/\theta)$.

3. The Advisor chooses whether or not to disclose her signal to the Leader in message $m \in \{s, \varnothing\}$.

4. The Leader observes message m and decides which policy $a \in \mathbf{R}$ to implement.

We denote the Leader's and the Advisor's preferences by $U_L(\cdot)$ and $U_A(\cdot)$ correspondingly. We assume that the Leader wants to match the state of the world (w) and suffers a disutility from her policy choice's departures from w:

$$u_L(a; m, \theta) = -(a - w)^2. \tag{1}$$

Note that we assume that the Leader does not derive utility from the Advisor's competence directly; the Advisor's competence affects the Leader's utility only indirectly, as it alters the informativeness of advice (see below).

The Advisor exhibits more conservative preference than the Leader and wishes the policy to match $c \cdot 0 + (1 - c) \cdot w$, where $c \in [0, 1]$ measures the Advisor's conservatism, with higher c corresponding to a more conservative preference. The Advisor suffers a disutility from the

[2] An alternative formulation of our formal framework where distortionary messages can permit separation between types is one in which the advisors can, in a continuous fashion, introduce "scrambling" into the message, and the Leader must invest into verification to be able to assess what, if any, part of the message truthfully communicates the signal received by the Advisor. In such a framework, the Advisor's strategy would again depend on her competence analogously to the present model, with less competent types choosing less distortion. We are grateful to Carlo Prato for a helpful conversation that clarified this point.

departures of the Leader's policy a from her own, more conservative, preference. Apart from depending on the Leader's policy choice, the Advisor's utility also depends on whether or not the signal she sends is informative. We assume, in particular, that every Advisor values office and receives a finite and positive office benefit Ψ if and only if she sends an informative signal (i.e., if $m \neq \varnothing$). (We can also interpret Ψ as measuring the Advisor's opportunity cost of maintaining her position as the Advisor to the current Leader. Lower Ψ means the outside options are more attractive and the value of continuing as the Leader's Advisor is less attractive.) Formally, the Advisor's utility is given by

$$u_A(m; s, \theta, c) = \begin{cases} -(a - (1-c) \cdot w)^2 & if\ m = \varnothing, \\ -(a - (1-c) \cdot w)^2 + \Psi & \text{else.} \end{cases} \tag{2}$$

The Advisor's competence θ shows how much weight the Leader should give her advice if advice is given. By way of interpretation, consider competence θ as a characteristic not of a single individual, but, instead, of a bureaucratic agency, whose information is constrained by formal and informal institutional rules. (An alternative interpretation of θ along these lines is as a budget allocated to the agency.) In this way, if the Advisor is dismissed, the same constraints will affect her successor. Because of that, it becomes sequentially rational for the Leader to replace the Advisor for failing to provide an informative message, or, equivalently, to commit to rewarding the Advisor who does send such a message.

The solution concept is weak Perfect Bayesian Equilibrium, which imposes the following requirements. Following the Advisor's message m, the Leader forms posterior belief $\mu(w|m, \theta)$ derived from the Advisor's strategy using Bayes's rule, where $\mu(\cdot)$ denotes the probability that the state of the world is w conditional on the message m the Leader observes and the competence of the Advisor θ. The Leader chooses $a^*(m, \theta)$ such that it maximizes his expected utility, given his posterior beliefs

$$a^*(m,\theta) = \arg\max_a \int u_L(a,w)d\mu(w|m,\theta). \tag{3}$$

Further, the Advisor's message $m^*(\theta)$ maximizes the Advisor's expected utility, given her (correct) expectation of the Leader's policy choice

$$m^*(\theta) = \arg\max_{m\in\{s,\varnothing\}} u_A(a^*(m,\theta)). \tag{4}$$

Analysis

Preliminaries: the Trade-off Between Competence and Advice

To provide a general intuition for the potential benefit to the Leader of having a less than maximally competent advisor, we begin by abstracting away from the exact functional forms of the utilities and signal distributions. Suppose that the Leader anticipates the Advisor of competence θ sends an informative signal with probability $r(\theta)$ and an uninformative signal with complementary probability $(1 - r(\theta))$. We denote the expected utility the Leader gets when he observes the informative signal $\alpha(\theta)$. When the Leader does not observes the informative signal, he gets expected utility $\beta(\theta) < \alpha(\theta)$.

The Leader's expected utility is, then,

$$E[u_L(\theta)] = U_L(\theta) = r(\theta)\cdot\alpha(\theta) + (1 - r(\theta))\cdot\beta(\theta).$$

Differentiating the Leader's utility with respect to θ, we get

$$\frac{\partial U_L(\theta)}{\partial\theta} = r'(\theta)\cdot(\alpha(\theta) - \beta(\theta)) + r(\theta)\cdot\alpha'(\theta) + (1 - r(\theta))\cdot\beta'(\theta). \tag{5}$$

Comparing this expression to 0, we obtain the following remark:

Remark 1. *The Leader's utility decreases in the Advisor's competence when*

$$r'(\theta) \cdot (\alpha(\theta) - \beta(\theta)) < -r(\theta) \cdot \alpha'(\theta) - (1 - r(\theta)) \cdot \beta'(\theta), \tag{6}$$

and increases in her competence otherwise.

In the inequality 6, $(\alpha(\theta) - \beta(\theta))$ represents the difference between the Leader's utility with and without information and shows the value of information to the Leader. When $\alpha(\theta)$ and $\beta(\theta)$ are close, information revelation does not benefit the leader. $r'(\theta)$ represents the change in the probability of acquiring information following a shift in the advisor's competence. Therefore, the left-hand side of inequality 6 shows the marginal loss (if $r'(\theta)$ is negative) or gain (if $r'(\theta)$ is positive) in information, weighted by its importance to the Leader.

The right-hand side of inequality 6 represents the Leader's marginal utility from acquiring information, assuming a fixed probability of receiving it. Remark 1, thus, holds that the Leader's utility decreases as the Advisor's competence increases when the marginal loss of information per se multiplied by its importance to the Leader, is lower than the marginal utility gained from acquiring information.

Of course, the question of whether this condition is consistent with a properly micro-founded agency model remains open. In the next section, we show when such condition is supported in the context of our model.

Equilibrium Analysis

To characterize the equilibrium behavior in our model, we proceed by backward induction.

When the Leader observes the (directly) informative message $m = s$, he adopts a policy equal to the mean of the posterior distribution conditional on the message he observes. When the Advisor does not convey her signal, i.e., $m = \varnothing$, the Leader chooses policy equal to the mean of his posterior belief conditional on the Advisor's having concealed her signal.

Note that both $m = s$ and $m = \emptyset$ may be informative in equilibrium; but the former is so intrinsically, whereas the informational content of the latter is determined by its use in equilibrium. Henceforth, we will use the term "informative message" solely to indicate the intrinsically informative message $m = s$. Given the Leader's anticipated policy response, the Advisor sends an informative message when her utility from revelation exceeds her utility from concealment. The following proposition characterizes the Leader's and the Advisor's equilibrium strategies:

Proposition 1. *In equilibrium, the Advisor sends an informative signal to the Leader if and only if*

1. *she is not too conservative,* $c \leq 1/2$*; or*

2. *she is sufficiently conservative,* $c \geq 1/2$ *and the signal is close enough to the mean state,* $-\hat{s}(\theta, \cdot) < s < \hat{s}(\theta, \cdot)$*, where* $\hat{s}(\theta, \cdot) \equiv \sqrt{\Psi} \cdot (1 + \frac{1}{\theta}) \cdot \frac{1}{\sqrt{2c-1}}$.

Otherwise, the Advisor conceals her information.

The Leader implements policy

$$a^* = \begin{cases} \frac{m}{1+1/\theta}, & \text{if } m = s; \\ 0, & \text{if } m = \varnothing. \end{cases} \tag{7}$$

Proof. See Appendix A. □

Note that the conditions for revelation in equilibrium are such that when the Advisor chooses $m = \varnothing$, the Leader's posterior belief is symmetric about the prior mean, though more dispersed than the prior. Figure 1 shows the thresholds $\pm\hat{s}(\theta, \cdot)$ as a function of the Advisor's competence θ assuming $c = 1$. The shaded area depicts signals that an Advisor of competence θ reveals to the Leader.

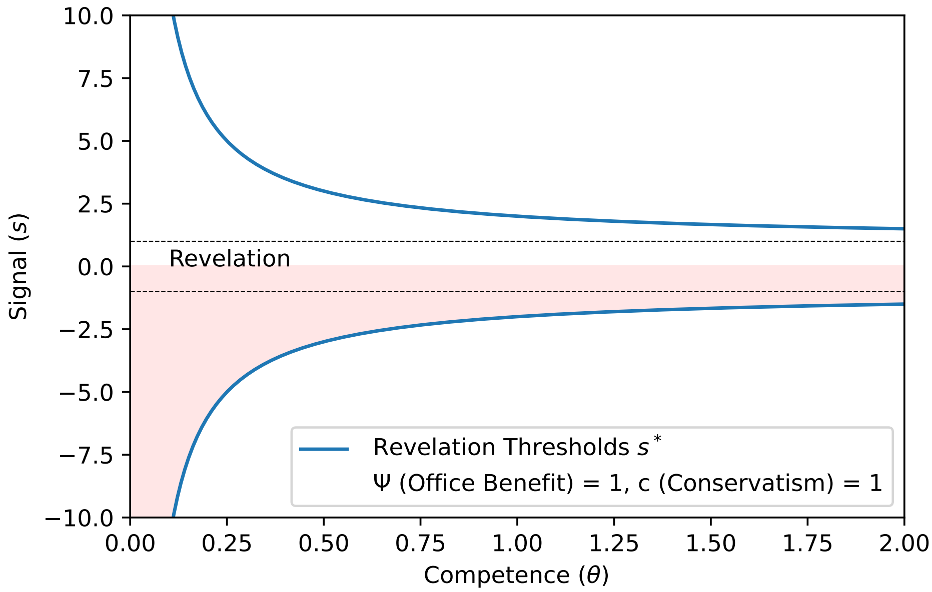

Figure 1: The solid curves show the signal thresholds $\pm\hat{s}(\theta)$ as a function of the Advisor's qualification θ when $c = 1$. When the Advisor receives a signal $s \in [-\hat{s}(\theta), \hat{s}(\theta)]$, she reveals her signal to the Leader. Dashed lines are asymptotes of the thresholds $\pm\hat{s}(\theta)$ when $\Psi = 1$.

As we can see in Figure 1, the higher the Advisor's competence, the smaller the range of informative messages the Advisor sends to the Leader. The next proposition characterizes how the Advisor's incentives to share information with the Leader, given the Leader's anticipated policy response, vary with the parameters of the model.

The final policy the Leader adopts depends on the information shared by the advisor and the advisor's competence, θ. The higher the Advisor's competence θ, the more the Leader relies on the Advisor's message when selecting policy. Thus, holding fixed s, the greater the Advisor's competence, the larger the difference between the policy the Leader selects after receiving the message and Advisor's ideal point, which, in turn, more strongly discourages the Advisor of higher θ from disclosing s in the first place. This observation highlights the main trade-off that determines the Leader's utility: more competent Advisors are less likely to reveal information, forcing the Leader to balance the quality of advice and its availability.

The office benefit Ψ directly enters the Advisor's utility, encouraging her to reveal more information to the Leader. The Advisor's conservatism c captures the Advisor's preference for the prior mean vs. the current state. The higher is the Advisor's conservatism, the further is her ideal point from the policy the Leader chooses after $m = s$, on average, thereby discouraging revelation. Figure 2 illustrates these relationships at a fixed value of c (left panel) and fixed value of Ψ (right panel).

The following proposition summarizes the preceding observations:

Proposition 2. *In equilibrium, the Advisor's incentive to send an informative signal to the Leader*

1. *decreases with her competence;*

2. *increases in the Advisor's valuation of office, Ψ; and*

3. *decreases in the Advisor's conservatism, c.*

Proof. Follows from Proposition 1 by taking derivatives with respect to θ, Ψ, and c. □

Figure 2: Equilibrium Advisor's strategy as a function of competence for different values of Ψ and c.

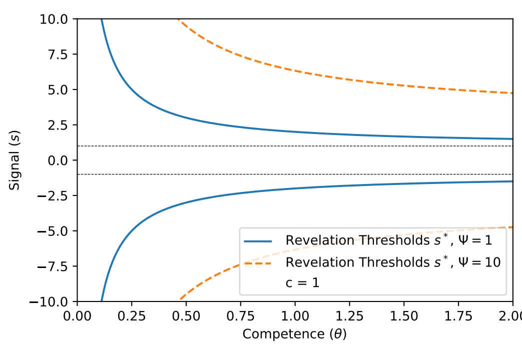

(a) The Advisor's revelation thresholds as a function of her competence θ at $c = 1$. Solid curves show the revelation thresholds for $\Psi = 1$. The dashed curves show the revelation thresholds for $\Psi = 10$.

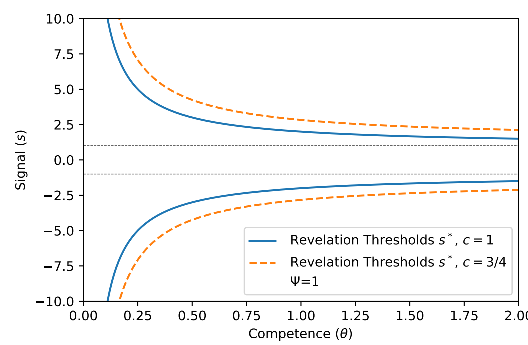

(b) The Advisor's revelation thresholds as a function of her competence θ at $\Psi = 1$. Solid curves show the revelation thresholds for $c = 1$. The dashed curves show the revelation thresholds for $c = 3/4$.

Returning to the trade-off between the quality and availability of advice described in Remark 1, the equilibrium probability that an Advisor of competence θ reveals her informative signal to the leader is $r(\theta) \equiv Pr[s \in \big(-\hat{s}(\theta), \hat{s}(\theta)\big)]$. The Leader's expected utility with an Advisor of competence θ is thus

$$E[u_L(\theta)] = U_L(\theta) = r(\theta) \cdot \frac{-1}{1+\theta} + (1 - r(\theta)) \cdot E[-w^2 | s \notin \big(-\hat{s}(\theta), \hat{s}(\theta)\big)]. \quad (8)$$

Recall that, conditional on not observing an informative message from the Advisor, the Leader infers that the signal the Advisor observed is not in $\big(-\hat{s}(\theta), \hat{s}(\theta)\big)$. For the closed form of the Leader's expected utility in equilibrium, see Appendix B.

The following proposition shows a key property of the Leader's expected utility as a function of Advisor competence, consistent with Remark 1:

Proposition 3. *There exists a unique threshold $\Psi^*(c)$ such that the Leader's equilibrium welfare is greatest with an Advisor of some finite competence for all $\Psi < \Psi^*(c)$, but with an Advisor of infinite competence otherwise.*

Proof. For proof see Appendix C. □

Proposition 3 shows that when the Advisors do not value office highly enough, the Leader is better off with an advisor of limited competence. While the Leader is always made better off by having better (more precise) information, advisors with high-quality information do not always deliver high-quality advice. Instead, as we show in Proposition 2, they are the most likely to conceal their knowledge to avoid significant policy changes. In addition to that, the more highly the Advisor values her position, the less likely she is to conceal information. Because higher types θ conceal more information than do lower types, in a setting with low office benefit (Ψ), the Leader is better off with an advisor of limited competence who reveals her lower quality information rather than with a more competent advisor who rarely provides her higher quality information.

Figure 3 shows the Leader's expected utility as a function of competence (θ) for different values of Ψ and c. Other things being equal, when the advisors value office highly (panel (a)), the Leader's expected utility increases in the Advisor's competence. However, as the office valuation (Ψ) decreases, the advisors of high competence begin to conceal more information from the Leader. Panel (b) illustrates the case in which the value of the position (Ψ) is low, and thus the Leader gets more useful advice (and higher expected utility) from a relatively low competence Advisor.

Figure 3: Expected Leader's utility as a function of the Advisor's competence for different Ψ

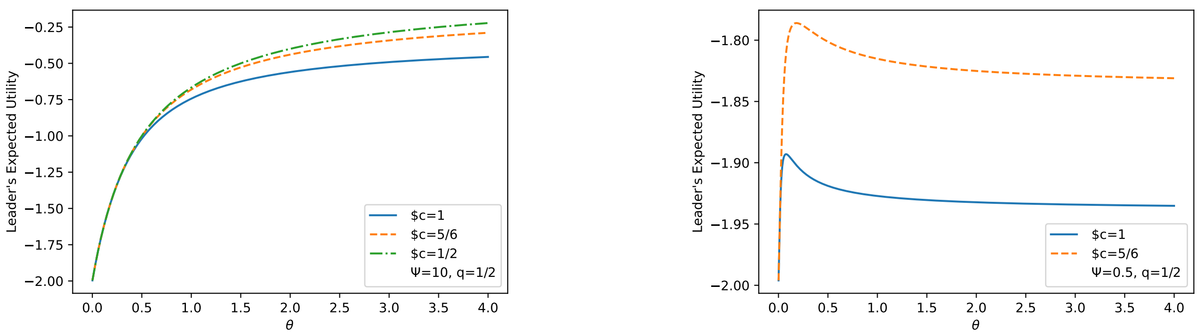

(a) The Leader's expected utility for different values of c, at $\Psi = 10$.

(b) The Leader's expected utility for different values of c, at $\Psi = 0.5$.

Note that the Advisor's valuation of office (Ψ) and her conservatism (c) essentially describe how co-aligned the Advisor is with the Leader. While the former directly measures the Advisor's incentives to remain in office with the current leadership, the latter affects the Advisor's valuation of the policies the Leader implements, affecting her revelation strategy and, thus, indirectly affecting the Leader. Our next result is consistent with the results from other models of communication: greater alignment of the Advisor's preferences with those of the Leader encourages revelation and benefits the Leader.

Proposition 4. *Both (a) the Leader's equilibrium utility holding fixed the Advisor's type, and (b) for $\Psi < \Psi^*(c)$, the type of Advisor θ that maximizes the Leader's equilibrium utility*

1. *increase in the Advisor's value of office, Ψ;*

2. *decrease in the Advisor's conservatism, c.*

Proof. See Appendix D for proof. □

We conclude by considering, in the context of our model, the possibility of bureaucrats' marginally "fudging" in communicating their evidence to political leaders. In the model we study, such a possibility is ruled out: the bureaucrats' options are to reveal their evidence as it stands or to withhold it altogether. Consider, though, the possibility that an Advisor can, without the Leader's detection, tilt the report in a preferred direction by some margin – perhaps because understanding the actual evidence requires complementary expertise, which is available to the bureaucrat but not to the political leader. How might such a possibility affect the expected equilibrium behavior? Assuming tilting is not detectable, we should expect the bureaucrats' reports to be tilted toward the status quo (perhaps by intervals dependent on the competence of the Advisor, if the latter is understood to affect how much tilting the Advisor can get away with without detection, or dependent on the extremity of the evidence). But then, when tilting the report is an option, the Leader should infer that the state is more extreme (in the same direction) than what the Advisor is reporting, and so choose the correspondingly more extreme policy. While the Advisor's and the Leader's choices will thus, nominally, change, the substantive aspects of the equilibrium play will remain intact.

Endogenous Information Acquisition

In this section, we consider an extension of our model that incorporates the Advisor's endogenous information acquisition. Maintaining the interpretation of the Advisor as a bureaucratic agency, we can think of the decision to run an information-generating policy experiment as the commissioning of a report, and the choice of the precision of the experiment as reflecting decisions about how much of the agency's resources to devote to the project.

Consider the following motivating example. The U.S. President wants to identify ways to save money by streamlining business operations in the Department of Defense, but does not know what specifically can be done or how much money could be saved by doing so. The department administration, which works closely with the defense industry and department insiders, responds by seeking the best possible information and so commissions a study from an independent advisory board and a highly regarded business consultancy, giving them full access to all relevant data throughout the department. While the department administration is seeking a most informative study, it also takes measures to ensure that, once this analysis of operations is acquired, it will control its revelation. (As it happens, the results of the study come as a shock, identifying extraordinary levels of wasteful spending and recommending a course of action projected to save a sum far in excess of what was expected, and the department takes steps to tighten the distribution of the study and its recommendations.)

The game we analyze below proceeds as follows. The Advisor, of competence θ, decides whether to conduct an experiment. If she does not, the Leader chooses a policy to implement, and the game ends. If she does, the Advisor chooses parameter $\tau \in (0, \theta]$ that characterizes the precision of the experiment and observes the experiment's realization: signal $s = w + \varepsilon$, where $\varepsilon \sim N(0, \tau)$. The Advisor, thus, chooses whether to run an experiment that could be relatively more or less informative, but cannot run a better experiment than her own quality permits.[3] We assume that the Leader does not observe τ, but his conjecture $\hat{\tau}$ must be correct in any pure-strategy equilibrium, consistent with the definition of equilibrium. Thus, although the Leader "knows" τ in equilibrium, the value of τ is not contractible. This assumption is particularly appropriate in the current setting, as agency bureaucracies have considerable discretion and relative insulation within their information-gathering and processing functions (see, e.g., the discussion of the above DoD study example in Whitlock and Woodward 2016; see also Nou, 2012; Roberts, 2009).

After the experiment (if the Advisor chose to run one), the Advisor chooses to reveal, or not, its outcome to the Leader. Once the revelation stage is complete, the Leader updates the policy based on the information he receives.

[3] We assume that there is no cost associated with information acquisition. This assumption is in line with our general focus on the indirect constraints on informative communication between the principal and the agent, in contrast to direct constraints such as the direct costs that the Advisor incurs in case of information acquisition. We relax the assumption of costless information in the appendix (see Appendix F).

In order to simplify the analysis in what follows, we set $c = 1$ – that is, we assume that the Advisor is maximally conservative, so that her most preferred policy is the status quo.

From Proposition 1, we have that the Advisor's only sequentially rational decision, after running an experiment, is to share her signal s with the Leader if and only if

$$-\sqrt{\Psi} \cdot (1 + \frac{1}{\tau}) < s < \sqrt{\Psi} \cdot (1 + \frac{1}{\tau}) \equiv \tilde{s}(\tau, \Psi). \tag{9}$$

Note that, during the revelation stage, the Advisor's competence θ does not explicitly affect the Advisor's revelation strategy, which entirely depends on the chosen experiment precision.

Prior to the Advisor's decision to run the experiment, the Advisor determines its precision, anticipating that she will reveal the signal if an only if (9) is satisfied. Two substantially different classes of equilibria might arise in the aftermath of this decision. In the first one, the Advisors of all competence levels choose null precision ($\tau^* = 0$), and the Leader always believes that the precision of the signal he observes must equate zero ($\hat{\tau} = 0$). In the second one, the Advisor chooses the highest possible precision available to him ($\tau^* = \theta$), and the Leader's beliefs are consistent with this strategy ($\hat{\tau} = \theta$). While both of them constitute weak Perfect Bayesian Equilibrium, there is a unique equilibrium that survives is consistent in the sense of Kreps and Wilson (1982). The following remark

Remark 2. *In the unique sequential equilibrium, the Advisor will choose the experiment of the highest possible precision, conditional on running the experiment.*[4]

Proof. See Appendix E.1. □

In effect, in equilibrium, the Advisor can never credibly commit to sabotaging the experi-

[4]Note that this equilibrium is the unique sequential equilibrium in the game, but one of two Bayesian Nash Equilibria. In the other Bayesian Nash equilibrium, the Leader always believes the experiment to be of the lowest possible precision ($\hat{\tau}^* = 0$) and the Advisors of all competence levels select null precision of the experiment ($\tau = 0$). This equilibrium, however, will not satisfy the belief consistency condition: Imagine the Leader's conjecture about the experiments' precision exceeds zero ($\hat{\tau}^n > 0$). Conditional on the Leader's updating, the Advisor must, then, choose maximal possible precision ($\tau^n(\hat{\tau}^n) = \theta$) as it guarantees the realization most centered around the Advisor's preferred policy outcome. Therefore, as $\hat{\tau}^n$ converges to zero, the Advisor's chosen precision will converge to θ instead of zero, violating the consistency condition.

ment. Holding the Leader's conjecture $\hat{\tau}$ fixed, increasing the precision of the experiment alters the signal realization distribution. The higher the precision, the less likely (in expectation) is the realization to fall further away from the Advisor's ideal point, and therefore, the Advisor will always have an incentive to deviate from any precision level that is lower than the maximum value, $\tau = \theta$. Therefore, $\hat{\tau} = \theta$ is the only conjecture that is consistent with equilibrium play, and so we proceed setting $\tau = \theta$ in the subsequent analysis.

Consider next the Advisor's decision to initiate the experiment, given she and the Leader behave sequentially rationally in all subsequent stages. The Advisor's expected utility from running the experiment is

$$
\begin{aligned}
&E[u_A(R=1)] = U_A(R=1) \\
&= Pr[s \in (-\tilde{s}(\tau=\theta,\Psi), \tilde{s}(\tau=\theta,\Psi)] \cdot \int_{-\hat{s}(\theta,\cdot)}^{\hat{s}(\theta,\cdot)} \left(-(\frac{x}{1+1/\theta})^2 + \Psi\right) \cdot \phi(\frac{x}{\sqrt{1+1/\theta}}) \cdot \sqrt{1+1/\theta}\, dx \\
&\qquad + Pr[s \in (-\tilde{s}(\tau=\theta,\Psi), \tilde{s}(\tau=\theta,\Psi)] \cdot 0, \qquad (10)
\end{aligned}
$$

where $\phi(x)$ is the probability density function of the standard normal distribution. When the Advisor chooses not run the experiment ($R = 0$), she gets utility zero. When she, instead, initiates the experiment, the Advisor must consider the Leader's reaction to every possible signal she might reveal given her chosen precision ($\tau = \theta$).

The Advisor's control over the information flow allows her to conceal realizations that lead to policy changes that are too drastic. For this reason, when facing a choice of whether to run an experiment, the Advisor always does so (assuming, of course, as we have been, that the experiment is costless to run).

Our next proposition summarizes the Advisor's equilibrium strategy and associated comparative statics.

Proposition 5. *When the information acquisition is endogenous,*

1. *the Advisor always runs the experiment; and*

2. *the Advisor reveals the result of the experiment when it falls within the interval* $[-\tilde{s}(\theta, \Psi), \tilde{s}(\theta, \Psi)]$ *and conceals if otherwise.*

Proof. See Appendix E.2. □

The analysis we present in this section reinforces the competence-revelation trade-off introduced in the baseline model. Conditional on having acquired information, an Advisor of higher competence continues to have a greater incentive to conceal this information. Given that the experiment realization remains private until the Advisor decides to share its realization, the Advisors of all types always acquire information and later conceal unfavorable, for them, outcomes.

Disclosure Requirement

In our model, the trade-off between the quality of advice and the likelihood of receiving it raises the possibility that the Leader is sometimes better off with a less competent advisor. Of course, the trade-off is a consequence of particular behavioral choices by the advisors, and the Leader has a preference over those: she prefers receiving advice to not receiving it and, furthermore, values receiving more rather than less competent advice. This raises the question of what institutional tools the Leader could use to induce behavioral choices she prefers to help mitigate the trade-off. In this section, we explore an institutional feature that requires the Advisor to reveal her information – the disclosure requirement – and study its equilibrium effects. We show that requiring disclosure can, indeed, mitigate the trade-off between the quality and the quantity of advice – but that that effect is fundamentally contingent on the nature of the Advisor's information. When the information is fixed, a disclosure requirement makes higher quality advisors more beneficial to the Leader. But when the advisor's information is endogenous, a disclosure requirement has the opposite effect – it makes higher quality advisors (weakly) less beneficial and lowers the Leader's expected utility.

To model the disclosure requirement, we modify the baseline environment to assume that the Leader immediately learns all information the Advisor observes, and the Advisor, then, always gets an office benefit Ψ. When the Advisor's information is fixed (exogenous), the disclosure requirement means that the Leader learns whatever the Advisor knows.[5]

Trivially, the disclosure requirement increases the value of Advisor competence, θ, to the Leader, ceteris paribus.

Remark 3. *Holding fixed the Advisor's information, the disclosure requirement increases the Leader's marginal utility from the Advisor competence.*

Thus, holding fixed the Advisor's information, a disclosure requirement fully eliminates the competence-information trade-off we saw in the baseline model. It forces the revelation of the information most valuable to the Leader, and it has the greatest impact precisely for the most competent advisors, who would, otherwise, have concealed the most information.

The left panel of Figure 4 contrasts the Leader's utility with no disclosure requirement and the Leader's utility with disclosure requirement, holding fixed the Advisor's information. When Advisor's information is exogenously given, the disclosure requirement always serves to improve Leader utility and information.

However, when we take into account the impact of the disclosure requirement on the Advisor's incentives to pursue information, the conclusion change. As in the previous section, to introduce endogenous information acquisition, we assume that the Advisor can decide whether to acquire information (to initiate the experiment), and if she does, she chooses its precision, τ. Under the disclosure requirement, then, if she chooses to acquire information, the realization of the experiment is immediately observed by the Leader. When the Advisor initiates the experiment, she receives office benefit Ψ, and she gets utility zero otherwise.

We begin our analysis with a remark that mirrors Remark 2 for the endogenous infor-

[5] We set aside the analysis of the possibility of explicitly endogenizing Ψ. That said, in effect, being able to acquire information affects Ψ endogenously, and the effect is different for different τ. This parallels the effect of the size of the discretionary window on the benefits of office for *zealots* in Gailmard and Patty (2007).

Figure 4: Expected Leader's utility as a function of the Advisor's competence for different Ψ

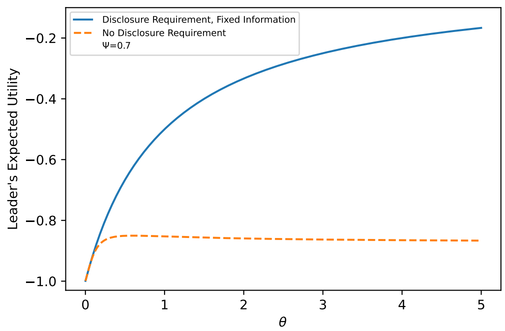

(a) The Leader's expected utility with and without disclosure requirement, given exogenous information acquisition.

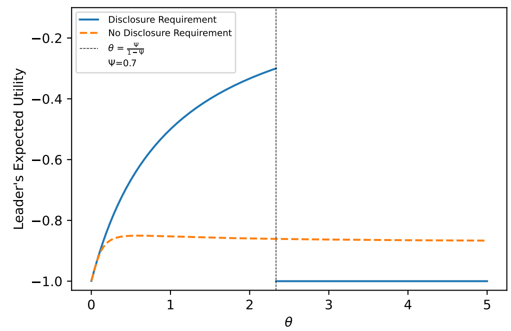

(b) The Leader's expected utility with and without disclosure requirement, given endogenous information acquisition.

mation extension of the baseline model:

Remark 4. *In the unique sequential equilibrium of the game with endogenous information acquisition and a disclosure requirement, the Advisor will choose the experiment of the highest possible precision, conditional on running the experiment.*

Proof. See Appendix G.1. □

Although all information acquired by the Advisor becomes immediately available to the Leader, the Advisor can never commit credibly to a precision that is less than the maximum. Conditional on running the experiment, higher precision guarantees that the experiment's realization ism, in expectation, closer to the Advisor's most preferred policy, and so, in equilibrium, the Advisor sets $\tau = \theta$.

When the Advisor proceeds with the experiment ($R = 1$), she gets utility

$$E[u_A(R = 1)] = \int_{-\infty}^{\infty} (-(\frac{x}{1 + 1/\theta})^2 + \Psi) \cdot \phi(\frac{x}{\sqrt{1 + 1/\theta}}) \cdot \sqrt{1 + 1/\theta}\, dx, \tag{11}$$

where $\phi(x)$ is the probability density function of the standard normal distribution. When the Advisor chooses not run the experiment ($R = 0$), she gets utility zero. Because the decision to run the experiment leads to policy change (because of the disclosure requirement), the

Advisor has to consider the Leader's reaction to every possible realization of the signal. Our next proposition summarizes the Advisor's equilibrium strategy and associated comparative statics.

Proposition 6. *When information acquisition is endogenous and there is a disclosure requirement,*

1. *the Advisor runs the experiment if her competence is below the threshold*

$$\hat{\theta}(\Psi) \equiv \begin{cases} \frac{\Psi}{1-\Psi} & if \Psi \leq 1 \\ \infty & if\ \Psi > 1 \end{cases} \tag{12}$$

and does not run it otherwise; and

2. *the threshold $\hat{\theta}(\Psi)$ increases in the Advisor's valuation of office Ψ.*

Proof. See Appendix G.2. □

Proposition 6 highlights the re-emergence of the equilibrium trade-off between competence and advice. Importantly, the disclosure requirement changes the set of equilibrium outcomes, discouraging more competent Advisors from acquiring information relative to settings with no disclosure requirement. A key intuition is that if the high-competence Advisor is not obligated to reveal, she may still choose to reveal when the signal, though ex ante unlikely, turns out to be close enough to the status quo. (Although, in our model, running the experiment does not have a direct cost for the Advisor, there is a range of costs that the Advisor would be willing to incur to run the experiment because the Advisor values the office benefit. This is true even if the disclosure of the information she acquires is relatively unlikely.) The disclosure requirement takes that discretion away from the Advisor, whose decision to run the experiment now turns on the ex ante distribution of signals. For sufficiently competent advisors, running the experiment under the disclosure requirement is a lottery with a lower expected value than the certain forfeiture of Ψ.

Note that the finding of the equilibrium impact of the disclosure requirement contrasts with the key result in Di Pei (2015). Di Pei extends the canonical cheap talk settings to allow for endogenous information acquisition where the sender can select the information structure prior to the communication stage. Under the assumption that the acquisition cost is increasing in information precision, the sender never obtains information she might then choose to conceal; instead, she lowers the precision, reducing acquisition costs. Thus, the game always resolves in full communication, and the disclosure requirement in Di Pei (2015) does not affect the sender's equilibrium play. Gentzkow and Kamenica (2017) establish a similar result in settings with verifiable messages. A key element that explains the difference between our predictions and those of Gentzkow and Kamenica and Di Pei is the extent of flexibility in the sender's choice of information structure; the expectation of full revelation in those papers depends on such flexibility. In contrast, in our model, because the Advisor lacks the ability to commit credibly to less-than-maximal information precision, she prefers to conceal certain obtainable information, and, therefore, enforced disclosure always discourages information acquisition.

As the analysis above demonstrates, the disclosure requirement does not necessarily serve to improve Leader utility when information acquisition is endogenous. Instead, it can backfire and result in the Leader receiving less information specifically when that information would be most valuable, i.e., when the Advisor is highly competent. Thus, while one might, ignoring the strategic effects, think that the setting with highly competent advisors is the setting in which the disclosure requirement would be most valuable to the leaders, this intuition is exactly wrong when information is endogenous: leaders should be less eager to require disclosure when the advising agency is highly competent, lest it destroy the advisor's willingness to acquire information.

Our next proposition generalizes the above intuition and characterizes conditions under which the Leader is better off with the disclosure requirement than without it.

Proposition 7. *When information acquisition is endogenous and disclosure is required,*

1. *Leader utility strictly increases in Advisor competence if and only if* $\theta < \frac{\Psi}{1-\Psi}$; *and*

2. *the Leader favors having a disclosure requirement when* $\theta < \frac{\Psi}{1-\Psi}$, *and prefers not having it otherwise.*

The right panel of figure 4 shows Leader utility with and without tge disclosure requirement, assuming endogenous information acquisition. When Advisor competence is sufficiently low, the disclosure requirement improves Leader utility. But if Advisor competence exceeds $\hat{\theta}(\Psi)$ threshold, the Advisor does not acquire information, and Leader utility is lower than it would be with no disclosure.

Discussion and Conclusion

We establish in general terms the existence of a trade-off between the quality of the information the advisor has to offer and her willingness to reveal it. The setting for this result – an advisor/agent with more conservative preferences than has the policy-maker, verifiable information, and limited rewards/punishments for the agent – is ubiquitous in modern, expertise-reliant governance. It is robust to endogenous information acquisition and is even exacerbated by endogenous information acquisition with sufficient publicity. These results suggest the need for a more nuanced approach in place of the common assumption that a policy-maker will fair better with, and thus prefer, a more competent advisor, and calls into question the presumption that the utilization of less competent advisors when more competent ones are available is a sign of corruption or dysfunction.

To be sure, especially in the context of advisors as bureaucratic agencies, advice-giving is but one of the functions that agents perform, and we abstract away from policy implementation, most notably. We bracket the latter function not because we believe it to be secondary, but to focus on a strategic mechanism that has received little prior attention. A subsequent analysis would do well to consider interactions between advising and implementing, which may uncover important relationships between the strategically contingent quality of advice and the quality of policy implementation based on such advice, but such an inquiry goes beyond the scope of this paper.

Our analysis of the disclosure requirement suggests that there is some, if limited, value to that institutional lever from the standpoint of improving the informational quality of leaders' choices. A different possibility worth considering from the same standpoint is multiple independent information sources – e.g., multiple, potentially competing, agencies. Such a possibility presents a new and complex set of strategic challenges. To see a key intuition, consider an advisor's expectation of when revealing her signal would reinforce or, alternatively, undermine the Leader's (intermediate) posterior based on the messages from other advisors. If another advisor's message is on the same side of the status quo, the value to an advisor of revealing her signal is, all else equal, negative. On the other hand, if another advisor's message is on the opposite side of the status quo, the value of revealing could become positive if it forces the Leader to update closer to the status quo. As this makes clear, advisors, then, have substantial incentives to collude in order to avoid sending mutually reinforcing messages. While a detailed examination of the effects of multiple advisors is beyond the scope of our analysis here, the above discussion also indicates why adding advisors is not a magic bullet for resolving the trade-off between the quality and the quantity of advice: when advisors can coordinate, the case in which the set of advisors are high-competence and their individual signals are far from the status quo (and so, likely to be on the same side of the status quo) is one in which the advisors' incentives against revelation are particularly strong and in which the Leader faces an especially difficult challenge in inducing revelation.

Another intuitive approach to mitigating the problem facing the Leader revolves around commitment in policy choice. The idea of such a commitment in the context of our model can be instructively compared to the commitment-based approach analyzed in the seminal work by Gilligan and Krehbiel (1989). Gilligan and Krehbiel show, in the context of an interaction between a legislative committee and the less informed median voter on the legislative floor, that closed rule, which limits the amendments that the floor can make to the committee's proposal before the up-or-down vote, can increase the committee's incentives to reveal its private information. The stark case of a perfectly conservative Advisor makes it clear, however, why a commitment of this kind cannot help the Leader in the setting we analyze above: Because the Advisor's preference is to maintain the status quo, and so is against revealing information to the Leader, the option of deferring to the sender's proposal is moot. But there is a different kind of commitment-based institution that could work for the Leader in our setting. If the Leader could be credibly committed to making only small changes to the status quo – in other words, if the Leader could be constrained to adopt a conservative posture, in effect, bringing his induced actions closer to the Advisor's most preferred actions – then the downside of revelation for the Advisor would decrease. This approach, however, entails substantial welfare losses for the Leader: notably, unlike in Gilligan and Krehbiel's setting, the Leader is unable to take advantage of the Advisor expertise to avoid big utility losses when the state of the world is a large departure from the prior (in our model, very high or very low). Thus, although the commitment approach may appear intuitively appealing here, it does not, at least on the basis of these very preliminary considerations, seem particularly promising from the standpoint of leader welfare. Of course, a more definitive conclusion requires a considerably more detailed analysis, which we leave to subsequent work.

Appendices

A Equilibrium Characterization

The Leader's expected utility from implementing a policy a after observing an informative message m will be

$$E[u_L(a;m)] = \int_{-\infty}^{\infty} -(x-a)^2 \cdot \frac{\frac{1}{2\cdot\pi} \cdot \frac{1}{\sqrt{1/\theta}} \cdot \frac{1}{\sqrt{1/1}} \cdot e^{-(x^2\cdot 1+(m-x)^2\cdot\theta)/2}}{\frac{1}{\sqrt{2\cdot\pi}} \cdot \frac{1}{\sqrt{1/1+1/\theta}} \cdot e^{-m^2/(1/1+1/\theta)/2}} dx, \tag{13}$$

Where the numerator of the fraction above represents the joint probability distribution of the state of the world (w) and random noise (ε), and the denominator shows the probability distribution of the message the advisor observes and shares with the Leader ($s = w + \varepsilon$).

The expected utility the Leader gets is, thus,

$$E[u_L(a;m)] = -\frac{1 + \theta + (a \cdot (1+\theta) - m \cdot \theta)^2}{(1+\theta)^2}. \tag{14}$$

To maximize his expected utility, the Leader who observes an informative message $m \neq \varnothing$ should select a policy $a^*(m \neq \varnothing) = \frac{m}{1+1/\theta}$. In the absence of information revelation, the Leader selects optimal policy denoted as $d^*(\theta)$.

The Advisor, then, sends an informative message when her expected utility from revealing this information exceeds her expected utility from concealing this information. [CORRECT BELOW FOR $C < 1/2$!]

$$\begin{aligned} E[u_A(m = s; s, \theta, c)] &= -(\frac{s}{1+1/\theta} - (1-c) \cdot \frac{s}{1+1/\theta})^2 + \Psi \\ > E[u_A(m = \varnothing; c, \theta)] &= -(d(\theta) - (1-c) \cdot \frac{s}{1+1/\theta})^2. \end{aligned} \tag{15}$$

Thus, she shares the signal with the Leader when she observes signal s

$$-\sqrt{\Psi + d(\theta)^2} \cdot (1 + \frac{1}{\theta}) \cdot \frac{1}{\sqrt{2c-1}} < s < \sqrt{\Psi + d(\theta)^2} \cdot (1 + \frac{1}{\theta}) \cdot \frac{1}{\sqrt{2c-1}} \tag{16}$$

and does not reveal her signal otherwise.

Finally, note that because the Advisor's strategy is symmetric around 0, in the absence of informative message, the Leader's optimal strategy is to implement policy $d^*(\theta) = 0$.

B Non-monotonicity of the Leader's Preference

Before we begin, let us note that the Advisor's conservatism (c) affects only the Advisor's incentives to reveal information to the Leader. Therefore, we can simplify the following analysis by denoting $\frac{\Psi}{2c-1}$ as new Ψ. In this case, all the results will held and comparative statics wrt c will mirror the inverse of comparative statics wrt Ψ.

The Leader's expected utility:

$$E[u_L(\theta)] = \underbrace{Pr[s \in (-\hat{s}, \hat{s})]}_{\substack{\text{Advisor sends} \\ \text{informative message}}} \cdot \underbrace{\frac{-1}{1+\theta}}_{\substack{\text{Leader's expected utility} \\ \text{after informative signal}}} + \underbrace{Pr[s \notin (-\hat{s}, \hat{s})]}_{\substack{\text{Advisor does not} \\ \text{send informative message}}} \cdot E[-w^2 | s \notin (-\hat{s}, \hat{s})], \tag{17}$$

where:

$$\begin{aligned} A &\equiv Pr[s \in [-\hat{s}, \hat{s}]] \times \frac{-1}{1+\theta} \\ &= \left(\Phi(\hat{s}/\sqrt{1+1/\theta}) - \Phi(-\hat{s}/\sqrt{1+1/\theta})\right) \times \frac{-1}{1+\theta}. \end{aligned} \tag{18}$$

and

$$
\begin{aligned}
B &\equiv Pr[s \notin (-\hat{s}, \hat{s})] \times E[-w^2 | s \notin [-\hat{s}, \hat{s}]] \\
&= Pr[s < -\hat{s}] \times E[-w^2 | s < -\hat{s}] + Pr[s > \hat{s}] \times E[-w^2 | s > \hat{s}] \\
&= Pr[s < -\hat{s}] \times (\int_{-\infty}^{\infty} \int_{-\infty}^{-\hat{s}-y} -x^2 \frac{f_{w,\varepsilon}(x,y)}{Pr[s < -\hat{s}]} dxdy) \\
&+ Pr[s > \hat{s}] \times (\int_{-\infty}^{\infty} \int_{\hat{s}-y}^{\infty} -x^2 \frac{f_{w,\varepsilon}(x,y)}{Pr[s > \hat{s}]}) dxdy) \\
&= (\int_{-\infty}^{\infty} \int_{-\infty}^{-\hat{s}-y} -x^2 \frac{1}{2\pi} \frac{1}{\sqrt{1/\theta}} e^{-\frac{1}{2}(x^2 + \frac{y^2}{1/\theta})} dxdy \\
&+ \int_{-\infty}^{\infty} \int_{\hat{s}-y}^{\infty} -x^2 \frac{1}{2\pi} \frac{1}{\sqrt{1/\theta}} e^{-\frac{1}{2}(x^2 + \frac{y^2}{1/\theta})} dxdy) \\
&= \frac{1}{2\pi} \frac{1}{\sqrt{1/\theta}} \int_{-\infty}^{\infty} -\sqrt{2\pi} \times e^{\frac{\theta y^2}{2}} \Big((\hat{s} - y)\phi\big(\hat{s} - y\big) \\
&+ (\hat{s} + y)\phi\big(\hat{s} + y\big) + 2 - \Phi(\hat{s} - y) - \Phi(\hat{s} + y) \Big) dy \\
&= \int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{1/\theta}} \times e^{-\theta y^2/2} \Big((\hat{s} - y)\phi\big(\hat{s} - y\big) \\
&+ (\hat{s} + y)\phi\big(\hat{s} + y\big) + 2 - \Phi(\hat{s} - y) - \Phi(\hat{s} + y) \Big) dy.
\end{aligned} \tag{19}
$$

Let us denote

$$
\begin{aligned}
g(a) \equiv \int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{1/\theta}} \times e^{-\theta y^2/2} \Big((\hat{s} - y)\phi\big(\hat{s} - y\big) + (\hat{s} + y)\phi\big(\hat{s} + y\big) \\
+ 2 - \Phi(\sqrt{a}(\hat{s} - y)) - \Phi(\sqrt{a}(\hat{s} + y)) \Big) dy.
\end{aligned} \tag{20}
$$

Note that $g(1) = B$ and our objective is to compute $g(1)$. Let us start by computing $g(0)$

$$
\begin{aligned}
g(0) = \int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{1/\theta}} e^{-\theta y^2/2} \Big((\hat{s} - y)\phi\big(\hat{s} - y\big) + (\hat{s} + y)\phi\big(\hat{s} + y\big) + 1 \Big) dy \\
= -1 - \frac{e^{\frac{-\hat{s}^2}{2(1+1/\theta)}} \sqrt{2/\pi}\hat{s}}{\sqrt{(1 + 1/\theta)^3}}.
\end{aligned} \tag{21}
$$

Now we compute derivative of $g(a)$ with respect to a. By Leibniz integral rule[6]

$$
\begin{aligned}
\frac{\partial g(a)}{\partial a} &= \frac{\partial}{\partial a}\int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}} \times e^{-\theta y^2/2}\big((\hat{s}-y)\phi(-(\hat{s}-y)) + (\hat{s}+y)\phi(-(\hat{s}+y)) \\
&\qquad +2 - \Phi(\sqrt{a}(\hat{s}-y)) - \Phi(\sqrt{a}(\hat{s}+y))\big)dy \\
&= \int_{-\infty}^{\infty} \frac{\partial}{\partial a}\Big(-\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}} \times e^{-\theta y^2/2}\big((\hat{s}-y)\phi(-(\hat{s}-y)) + (\hat{s}+y)\phi(-(\hat{s}+y)) \\
&\qquad +2 - \Phi(\sqrt{a}(\hat{s}-y)) - \Phi(\sqrt{a}(\hat{s}+y))\big)\Big)dy \\
&= \int_{-\infty}^{\infty} -\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}}e^{-\frac{\theta y^2}{2}}\Big(-\phi(\sqrt{a}(\hat{s}-y))\frac{1}{2\sqrt{a}}(\hat{s}-y) - \phi(\sqrt{a}(\hat{s}+y))\frac{1}{2\sqrt{a}}(\hat{s}+y)\Big)dy \\
&= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{1/\theta}}e^{-\frac{\theta y^2}{2}}\frac{1}{2\sqrt{a}}\Big(\frac{e^{-\frac{a(\hat{s}-y)^2}{2}}}{\sqrt{2\pi}}(\hat{s}-y) + \frac{e^{-\frac{a(\hat{s}+y)^2}{2}}}{\sqrt{2\pi}}(\hat{s}+y)\Big)dy \\
&= \int_{-\infty}^{\infty} e^{-\frac{\theta y^2}{2}}\sqrt{\theta} \times \frac{e^{-\frac{a(\hat{s}-y)^2}{2}}(\hat{s}-y) + e^{-\frac{a(\hat{s}+y)^2}{2}}(\hat{s}+y)}{4\pi\sqrt{a}}dy \\
&= \hat{s} \times \frac{e^{-\frac{a\hat{s}^2\theta}{2(1+\frac{a}{\theta})}}}{\sqrt{2\pi}\sqrt{a}(1+\frac{a}{\theta})^{3/2}}.
\end{aligned} \tag{22}
$$

We now take an integral wrt a of $\frac{dg(a)}{da}$:

$$
\begin{aligned}
g(a) &= \int \hat{s} \times \frac{e^{-\frac{a\hat{s}^2\theta}{2(a+\theta)}}}{\sqrt{2\pi}\sqrt{a}(1+\frac{a}{\theta})^{3/2}}da \\
&= 2\Phi\Big(\frac{\hat{s}}{\sqrt{\frac{1}{a}+\frac{1}{\theta}}}\Big) - 1 + C,
\end{aligned} \tag{23}
$$

where C is unknown constant. Finally, let us note that

$$
g(a=0) = -1 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}} \tag{24}
$$

and

$$
g(a=0) = \lim_{a\to 0} 2\cdot\Phi\Big(\frac{\hat{s}}{\sqrt{\frac{1}{a}+\frac{1}{\theta}}}\Big) - 1 + C = C, \tag{25}
$$

where equation 24 is corollary of equation 21 and equation 25 is corollary of equation 23.

[6] Leibniz integral rule applies because the integral of partial derivative converges link

Therefore

$$C = -1 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}}. \tag{26}$$

Finally,

$$\begin{aligned} B &= \frac{Pr[s \notin [-\hat{s},\hat{s}]]}{Pr[s < -\hat{s}]} \times g(a=1) \\ &= \lim_{a\to 1} 2\cdot\Phi(\frac{\hat{s}}{\sqrt{\frac{1}{a}+\frac{1}{\theta}}}) - 2 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}} \\ &= (2\cdot\Phi(\frac{\hat{s}}{\sqrt{1+\frac{1}{\theta}}}) - 2 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}}) \end{aligned} \tag{27}$$

Therefore, the expected Leader's utility is

$$\begin{aligned} E[U_L(\theta)] = &\big(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) - 1\big)\cdot\frac{-1}{1+\theta} \\ &+(2\Phi(\frac{\hat{s}}{\sqrt{1+\frac{1}{\theta}}}) - 2 - \frac{e^{\frac{-\hat{s}^2\theta}{2(1+\theta)}}\sqrt{2/\pi}\hat{s}}{(1+1/\theta)^{3/2}}). \end{aligned} \tag{28}$$

C Existence of the finite optimum

The Leader's utility derivative wrt θ is:

$$\begin{aligned} \frac{\partial E[u_L(\theta)]}{\partial\theta} = \frac{1}{2(1+\theta)^2}\Big(&- \frac{\sqrt{2}e^{-\frac{\Psi(1+1/\theta)}{2}}\sqrt{\frac{\Psi\theta(1+\theta)}{\pi}}(2\theta+\Psi(1+\theta))}{\theta^2} \\ &+ \underbrace{2(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) - 1)}_{>0}\Big). \end{aligned} \tag{29}$$

Note that $\frac{\partial E[u_L(\theta)]}{\partial\theta}$ converges to 1 as θ approaches 0 and converges to ± 0 as θ approaches infinity, where the sign depends on whether

$$F(\Psi) \equiv 4 - e^{-\frac{\Psi}{2}}\sqrt{\frac{2}{\pi}}\sqrt{\Psi}(2+\Psi) - 4\Phi(\sqrt{\Psi}))$$

is positive or negative. $F(\Psi = 0) = 0$ and it decreases in Ψ for $\Psi < 1$ and increases in Ψ when $\Psi > 1$. Therefore, for $\Psi \in (0, 1)$, $\frac{\partial E[u_L(\theta)]}{\partial \theta}$ is negative at $\theta = \infty$ and, because $E[u_L(\theta)]$ is continuous and differentiable, the interior optimum will exist (note that $\Psi < 1$ is sufficient but not necessary condition for the existence of the interior optimum).

Let us denote the derivative of the Leader's utility with respect to the Advisor's qualification θ as $D(\theta, \Psi) \equiv \frac{\partial E[u_L|\theta]}{\partial \theta}$. Note that the Leader's utility reaches local maximum at $\hat{\theta}$ s.t. $D(\theta, \Psi) = 0$. By the implicit function theorem, in order to seek how Ψ affects the Leader's choice of interior optimal Advisor's qualification, one needs to compute

$$\frac{d\hat{\theta}(\Psi)}{d\Psi} = -\frac{\partial_\Psi D(\theta, \Psi)}{\partial_\theta D(\theta, \Psi)}. \tag{30}$$

Because we are looking for θ that maximizes the Leader's utility, $\partial_\theta D(\theta, \Psi)$ should not exceed zero. Therefore, sign of equation 30 mirrors sign of $\partial_\Psi D(\theta, \Psi)$. Because

$$\partial_\Psi D(\theta, \Psi) = \underbrace{\frac{e^{-\frac{\Psi\theta(1+\theta)}{2\theta}} \times \Psi}{2 \times \theta^2 \times \sqrt{2\pi} \times \sqrt{\Psi\theta(1+\theta)}}}_{>0} \times (\Psi(1+\theta) - \theta), \tag{31}$$

sign of $(\Psi(1+\theta)-\theta)$ determines whether $\hat{\theta}$ increases or decreases in Ψ. When $(\Psi(1+\theta)-\theta)$ is positive, optimal interior competence increases in Ψ, and it decreases in Ψ when $(\Psi(1+\theta)-\theta)$ is negative.

Note that $D(\theta, \Psi)$ reaches minimum at $\Psi = \frac{\theta}{1+\theta}$. Next, because $D(\theta = 0, \Psi) = \frac{1}{2}$ is positive while $D(\theta = \frac{\Psi}{1-\Psi}, \Psi) = -\frac{(1-\Psi)^2(3\sqrt{2}+\sqrt{e\cdot\pi}\cdot 2\cdot(1-2\cdot\Phi(1))}{2\cdot\sqrt{e\cdot\pi}}$ is negative, when $\Psi > \frac{\theta}{1+\theta}$, there will be interior maximum of the Leader's utility $\hat{\theta}$ s.t. $\hat{\theta} < \frac{\Psi}{1-\Psi}$. Because $\Psi > \frac{\theta}{1+\theta}$, this interior maximum $(\hat{\theta})$ increases in Ψ.

Finally, we prove that if the expected Leader's utility has a local maximum, this local maximum is unique. Once we prove this statement, we prove that **any** interior maximum increases in Ψ. First, note that derivative of the Leader's utility wrt θ is:

$$\frac{\partial E[u_L(\theta)]}{\partial\theta} = \underbrace{\frac{1}{2(1+\theta)^2}}_{>0}\Big(-\frac{\sqrt{2}e^{-\frac{\Psi(1+1/\theta)}{2}}\sqrt{\frac{\Psi\theta(1+\theta)}{\pi}}(2\theta+\Psi(1+\theta))}{\theta^2} \\ +2(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})-1)\Big). \tag{32}$$

Sign of $\frac{\partial E[u_L(\theta)]}{\partial\theta}$ mirrors the sign of

$$Interior(\theta) \equiv -\frac{\sqrt{2}e^{-\frac{\Psi(1+1/\theta)}{2}}\sqrt{\frac{\Psi\theta(1+\theta)}{\pi}}(2\theta+\Psi(1+\theta))}{\theta^2} + 2(2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})-1).$$

Now note that

$$\frac{\partial Interior(\theta)}{\partial\theta} = \frac{e^{-\frac{\Psi(1+1/\theta)}{2}}\sqrt{\frac{\Psi\theta(1+\theta)}{2\pi}}\Psi(\theta-\Psi(1+\theta)}{\theta^4}. \tag{33}$$

Therefore, $Interior(\theta)$ decreases in θ for $\theta < \frac{\Psi}{1-\Psi}$ and increases in θ for $\theta > \frac{\Psi}{1-\Psi}$. It implies that the expected Leader's utility can have no more than one local maximum.

As we just proved, the interior maximum of the Leader's utility exists if and only if $\Psi > \frac{\theta}{1+\theta}$. Therefore, when the interior maximum exists,

$$\partial_\Psi D(\theta,\Psi) = \frac{e^{-\frac{\Psi\theta(1+\theta)}{2\theta}} \times \Psi}{2\times\theta^2\times\sqrt{2\pi}\times\sqrt{\Psi\cdot\theta\cdot(1+\theta)}} \times (\Psi(1+\theta)-\theta) > 0.$$

Therefore, for any c there exists a unique threshold $\Psi^*(c)$ s.t. the Leader prefers interior competence over infinite competence for every $\Psi < \Psi^*(c)$.

D Comparative Statics

To prove part b of Proposition 4, note that as we show in Appendix C, the type of the Advisor θ that maximizes the Leader's equilibrium utility increases in the Advisor's value of office. Because, as we mention in Appendix B, we can denote $\frac{\Psi}{2c-1}$ as new Ψ to simplify analysis, it immediately follows that the type of the Advisor θ that maximizes the Leader's equilibrium utility decreases in the Advisor's conservatism.

E Experiment

E.1 Credible Precision

Assume that the Leader's beliefs of the precision the Advisor chooses are fixed at $\hat{\tau}$. Then, the Advisor's utility from running the private experiment is

$$\begin{aligned} E[u_A(R=1)] = (\Phi(\sqrt{\Psi} \cdot \frac{1+1/\hat{\tau}}{\sqrt{1+1/\tau}}) - \Phi(-\sqrt{\Psi} \cdot \frac{1+1/\hat{\tau}}{\sqrt{1+1/\tau}}))) \\ \cdot \Big(\int_{-\hat{s}(\hat{\tau},\cdot)}^{\hat{s}(\hat{\tau},\cdot)} -(\frac{x}{1+1/\hat{\tau}})^2 \cdot \phi(\frac{x}{\sqrt{1+1/\tau}}) \cdot \frac{1}{\sqrt{1+1/\tau}} \, dx + \Psi \Big). \end{aligned} \tag{34}$$

The Advisor's utility conditional on revelation always increases in τ : The higher precision experiment is (1) less likely to generate an outcome further away from the most preferred Advisor's outcome; and (2) more likely to result in revelation, allowing the Advisor to reap office benefit Ψ.

E.2 Information Acquisition

When the advisor acquires information, she gets utility

$$E[u_A(R=1)] = \frac{\phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) \cdot 2 \cdot \sqrt{\Psi} + \frac{(\Psi \cdot (1+1/\theta)-1) \cdot (2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})-1)}{\sqrt{(1+1/\theta)}}}{\sqrt{1+1/\theta}}. \tag{35}$$

Its derivative wrt θ is

$$\frac{\partial E[u_A(R=1)]}{\partial \theta} = \frac{2\phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) \cdot \frac{\hat{s}}{\sqrt{1+1/\theta}} - (2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) - 1)}{(1+\theta)^2}, \tag{36}$$

and the derivative of 36 wrt Ψ is equal to

$$-\frac{1}{(1+\theta)^2}\frac{1}{\Psi}\cdot\phi(\frac{\hat{s}}{\sqrt{1+1/\theta}})\cdot\sqrt{\pi}\cdot(\frac{\hat{s}}{\sqrt{1+1/\theta}})^3<0. \tag{37}$$

Therefore, $\frac{\partial E[u_A(R=1)]}{\partial\theta}$ decreases in Ψ and reaches maximum when Ψ converges to zero. Thus, the maximum possible $\frac{\partial E[u_A(R=1)]}{\partial\theta}$ is

$$\lim_{\Psi\to 0}\frac{\partial E[u_A(R=1)]}{\partial\theta}=0. \tag{38}$$

Thus, $E[u_A(R=1)]$ decreases in θ and reaches minimum when θ converges to infinity, where

$$\lim_{\theta\to+\infty}E[u_A(R=1)]=0. \tag{39}$$

Then the Advisor always runs the experiment.

F Costly Information Acquisition

In the baseline model we assume that the advisor does not Let us assume that information acquisition is costly and the Advisor incurs an expense of E is she decides to acquire information. The advisor gets utility 0 when she does not acquire information and utility $E[u_A(R=1)]-E$ when she does. Note that because $E[u_A(R=1)]$ decreases in θ (see Appendix E.2), for any E there will exist a unique threshold $\tilde{\theta}(E)$ such that the Advisor will always acquire information when her competence θ is such that $\theta\leq\tilde{\theta}(E)$ and will never acquire information when $\theta>\tilde{\theta}(E)$. Therefore, when information acquisition is costly, only Advisor's of sufficiently low competence run the experiment.

Furthermore, note that the advisor's utility following the information acquisition increases in the office benefit Ψ

$$\frac{\partial E[u_A(R=1)]}{\partial \Psi} = 2\Phi(\frac{\hat{s}}{\sqrt{1+1/\theta}}) - 1 > 0. \tag{40}$$

Therefore, for any expense that the advisor incurs when she acquires information (E), there exists a unique threshold $\tilde{\Psi}(E)$ such that the advisor acquires information if and only if the office benefit Ψ exceeds the threshold $\tilde{\Psi}(E)$.

The expense the Advisor incur when she acquires information will not affect her revelation strategy conditional on information acquisition. Following information acquisition, the Advisor reveals the result of the experiment when it falls within the interval $[-\tilde{s}(\theta, \Psi), \tilde{s}(\theta, \Psi)]$ and conceals information from the Leader otherwise.

G Disclosure Requirement

G.1 Credible Precision

The Advisor's utility from running the experiment when disclosure requirement is enforced is

$$E[u_A(R=1 \mid disclosure)] = \int_{-\infty}^{\infty} -(\frac{x}{1+1/\hat{\tau}})^2 \cdot \phi(\frac{x}{\sqrt{1+1/\tau}}) \cdot \frac{1}{\sqrt{1+1/\tau}} \, dx + \Psi, \tag{41}$$

which always increases in τ because the higher precision experiment is less likely to generate an outcome further away from the most preferred Advisor's outcome.

G.2 Information Acquisition with Disclosure Requirement

When the Advisor decides to acquire information, she gets utility

$$-\frac{1}{1+\theta} + \Psi, \tag{42}$$

and she gets utility zero otherwise. Therefore, the Advisor initiates information acquisition when

$$\theta < \begin{cases} \infty, & \Psi > 1, \\ \frac{\Psi}{1-\Psi}, & \Psi \leq 1. \end{cases} \tag{43}$$

References

Austen-Smith, David. 1990. "Information transmission in debate." *American Journal of political science* pp. 124–152.

Austen-Smith, David. 1993. "Interested experts and policy advice: Multiple referrals under open rule." *Games and Economic Behavior* 5(1):3–43.

Aybas, Yunus, C. and Steven Callander. 2023. "Efficient Cheap Talk in Complex Environments." *Working paper* .

Battaglini, Marco. 2002. "Multiple referrals and multidimensional cheap talk." *Econometrica* 70(4):1379–1401.

Bhattacharya, Sourav and Arijit Mukherjee. 2013. "Strategic information revelation when experts compete to influence." *The RAND Journal of Economics* 44(3):522–544.

Calvert, Randall L. 1985. "The value of biased information: A rational choice model of political advice." *The Journal of Politics* 47(2):530–555.

Canes-Wrone, Brandice, Michael C Herron and Kenneth W Shotts. 2001. "Leadership and pandering: A theory of executive policymaking." *American Journal of Political Science* pp. 532–550.

Che, Yeon-Koo and Navin Kartik. 2009. "Opinions as incentives." *Journal of Political Economy* 117(5):815–860.

Di Pei, Harry. 2015. "Communication with endogenous information acquisition." *Journal of Economic Theory* 160:132–149.

Dziuda, Wioletta. 2011. "Strategic argumentation." *Journal of Economic Theory* 146(4):1362–1397.

Fox, Justin. 2007. "Government transparency and policymaking." *Public choice* 131(1-2):23–44.

Gailmard, Sean and John W Patty. 2007. "Slackers and zealots: Civil service, policy discretion, and bureaucratic expertise." *American Journal of Political Science* 51(4):873–889.

Gailmard, Sean and John W Patty. 2012. "Formal models of bureaucracy." *Annual Review of Political Science* 15(1):353–377.

Gailmard, Sean and John W Patty. 2013. "Stovepiping." *Journal of Theoretical Politics* 25(3):388–411.

Gentzkow, Matthew and Emir Kamenica. 2017. "Disclosure of endogenous information." *Economic Theory Bulletin* 5(1):47–56.

Gilligan, Thomas W and Keith Krehbiel. 1989. "Asymmetric information and legislative rules with a heterogeneous committee." *American journal of political science* pp. 459–490.

Huq, Aziz and Tom Ginsburg. 2018. "How to lose a constitutional democracy." *UCLA L. Rev.* 65:78.

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101(6):2590–2615.

Kreps, David M. and Robert Wilson. 1982. "Sequential Equilibria." *Econometrica* 50(4):863–894.

Milgrom, Paul. 1981. "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics* 12(2):380–391.

Milgrom, Paul. 2008. "What the seller won't tell you: Persuasion and disclosure in markets." *Journal of Economic Perspectives* 22(2):115–131.

Nou, Jennifer. 2012. "Agency self-insulation under presidential review." *Harv. L. Rev.* 126:1755.

Ottaviani, Marco and Peter Norman Sørensen. 2006. "Reputational cheap talk." *The Rand journal of economics* 37(1):155–175.

Patty, John W. 2009. "The politics of biased information." *The Journal of Politics* 71(2):385–397.

Patty, John W and Ian R Turner. 2021. "Ex Post Review and Expert Policy Making: When Does Oversight Reduce Accountability?" *The Journal of Politics* 83(1):23–39.

Roberts, Patrick S. 2009. "How security agencies control change: Executive power and the quest for autonomy in the FBI and CIA." *Public Organization Review* 9(2):169–198.

Shin, Hyun Song. 1994. "The burden of proof in a game of persuasion." *Journal of Economic Theory* 64(1):253–264.

Sobel, Joel. 1993. "Information control in the principal-agent problem." *International Economic Review* pp. 259–269.

Sobel, Joel. 2013. "Giving and receiving advice." *Advances in economics and econometrics* 1:305–341.

Turner, Ian R. 2022. "Reviewing Procedure versus Judging Substance: How Increasing Bureaucratic Oversight Can Reduce Bureaucratic Accountability." *Journal of Political Institutions and Political Economy* 2(4):569–596.

Whitlock, Craig and Bob Woodward. 2016. "Pentagon buries evidence of $125 billion in bureaucratic waste." *The Washington Post, Accessed December 6th* .

Wolinsky, Asher. 2003. "Information transmission when the sender's preferences are uncertain." *Games and Economic Behavior* 42(2):319–326.